

RESIDUAL MINIMIZING MODEL INTERPLATION FOR PARAMETERIZED NONLINEAR DYNAMICAL SYSTEMS

PAUL G. CONSTANTINE* AND QIQI WANG†

Abstract. We present a method for approximating the solution of a parameterized, nonlinear dynamical system using an affine combination of solutions computed at other points in the input parameter space. The coefficients of the affine combination are computed with a nonlinear least squares procedure that minimizes the residual of the governing equations. The approximation properties of this residual minimizing scheme are comparable to existing reduced basis and POD-Galerkin model reduction methods, but its implementation requires only independent evaluations of the nonlinear forcing function. It is particularly appropriate when one wishes to approximate the states at a few points in time without time marching from the initial conditions. We prove some interesting characteristics of the scheme including an interpolatory property, and we present heuristics for mitigating the effects of the ill-conditioning and reducing the overall cost of the method. We apply the method to representative numerical examples from kinetics – a three state system with one parameter controlling the stiffness – and conductive heat transfer – a nonlinear parabolic PDE with a random field model for the thermal conductivity.

Key words. nonlinear dynamical systems, nonlinear equations, parameterized models, reduced order models, interpolation

1. Introduction. As computational capabilities grow, engineers and decision makers increasingly rely on simulation to aid in design and decision-making processes. However, the complexity of the models has kept pace with the growth in computing power, which has resulted in expensive computer models with complicated parametric dependence. For given parameter values, each costly evaluation of the model can require extensive time on massively parallel, high performance systems. Thus, exhaustive parameter studies exploring the relationships between input parameters and model outputs become infeasible; cheaper reduced order models are needed for sensitivity/uncertainty analysis, design optimization, and model calibration.

Reduced order modeling has become a very active field of research. Methods based on the Proper Orthogonal Decomposition (POD) have been shown to dramatically reduce the computational complexity for approximating the solution of linear dynamical systems or parameterized linear steady state problems [1, 22, 21, 5]. The success of such methods for linear models has spurred a slew of recent work on model reduction for nonlinear models [6, 7, 13, 10, 19]. In this vein, we focus on nonlinear dynamical systems that depend on a set of input parameters, but the method we develop can be modified for steady and/or linear parameterized models as well. The parameters may affect material properties, boundary conditions, forcing terms and/or model uncertainties; we do not consider parameterized initial conditions. Suppose we can afford to compute a full model solution at a few points in the parameter space, or equivalently, suppose we have access to a database of previously computed runs. How can we use those stored runs to cheaply estimate the model output at untested input parameter values?

In this paper, we propose an interpolation¹ method that employs an affine com-

*Department of Mechanical Engineering, Stanford University, Stanford, California 94305 (paul.constantine@stanford.edu).

†Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139 (qiqi@mit.edu).

¹In the work of [2, 7, 18], interpolation occurs in the spatial domain; our approximation interpolates in the parameter domain.

bination of the stored model evaluations and a nonlinear least squares procedure for computing the coefficients of the affine combination such that the equation residual is minimized. If one is interested in approximating the state vector at a new parameter value at only a few points in time, then constructing and solving the least squares problem will be significantly cheaper than solving the true model; this justifies the comparison to reduced order modeling. However, if one wishes to approximate the full time history of a few elements of the state vector, then the least squares problem may be more expensive to solve than the true model. We will highlight the former situation in the second numerical example.

Like standard ODE solvers, the method only requires evaluations of the forcing function from the dynamical system. Thus, implementation is straightforward using existing codes. However, unlike ODE solvers, the function evaluations can be performed independently to take advantage of parallel architectures. The model interpolation scheme itself has many appealing features including point-wise optimality by construction, reuse of stored model evaluations, and a strategy for adaptively choosing points in the parameter space for additional full model evaluations. We observe that the error in approximation decays like the eigenvalues of a covariance-like operator of the process as more model evaluations added. But we show that achieving a small residual requires the solution of an ill-conditioned least squares problem; we present a heuristic for taming the potentially unwieldy condition number.

The linear model of the data is a common feature of many model reduction and interpolation schemes including reduced basis methods [21, 26, 11, 5], kriging surfaces/Gaussian process emulators [23, 8, 12], and polynomial interpolation schemes [3, 28]. Unlike the kriging and polynomial schemes, our process utilizes the governing equations to construct the minimization problem that yields the coefficients, which invariably produces a more accurate interpolant; this is comparable to the reduced basis methods and POD-based model reduction techniques [7, 1]. However, in contrast to reduced basis methods and POD-based techniques, we formulate the least squares problem using only the equation residual, which requires minimal modifications to the full system solvers. Additionally, this residual-based formulation applies directly to nonlinear models, which have posed a persistent challenge for schemes based on Galerkin projection.

In Section 2, we pose the model problem and derive the residual minimizing scheme, including some specific details of the nonlinear least squares solver. We briefly show in Section 3 how kriging interpolation relates to the least squares procedure. In Section 4, we prove some properties of the residual minimizing scheme, including a lower bound on the average error, an interpolatory property, and some statements concerning convergence. Section 5 presents heuristics for adding full model evaluations, reducing the cost of the scheme, and managing the ill-conditioning in the least squares problems. In Section 6, we perform two numerical studies: (i) a simple nonlinear dynamical system with three state variables and one parameter controlling the stiffness, and (ii) a two-dimensional nonlinear parabolic PDE model of heat transfer with a random field model for thermal conductivity. Finally we conclude in Section 7 with a summary and directions for future work.

2. Residual Minimizing Model Interpolation. Let $x(t) = x(t, s)$ be a \mathbb{R}^p -valued process that satisfies the dynamical system

$$x' = f(x, t, s), \quad t \in [0, T], \quad s \in \mathcal{S}, \quad (2.1)$$

with initial condition $x(0) = x_0$. The space $\mathcal{S} \subset \mathbb{R}^d$ is the input parameter space, and we assume the process is bounded for all $s \in \mathcal{S}$ and $t \in [0, T]$. Note that the parameters affect only the dynamics of the system – not the initial condition. The \mathbb{R}^p -valued function f is nonlinear in the states x , and it often represents a discretized differential operator with forcing and boundary terms included appropriately. We assume that f is Lipschitz continuous and its Jacobian with respect to the states is nonsingular, which excludes systems with bifurcations. The process satisfying (2.1) is unique in the following sense: Let $y(t)$ be a differentiable \mathbb{R}^p -valued process with $y(0) = x_0$, and define the integrated residual $\phi(y)$ as

$$\phi(y) = \int_0^T \|y' - f(y, t, s)\|^2 dt, \quad (2.2)$$

where $\|\cdot\|$ is the standard Euclidean norm. If $\phi(y) = 0$, then clearly $y(t) = x(t)$ for $t \in [0, T]$. This residual-based definition of uniqueness underlies the model interpolation method. Given a discretization of the time domain $0 \leq t_1 < \dots < t_m \leq T$, we define the discretized residual $\phi_m(y)$ as

$$\phi(y) \approx \phi_m(y) = \sum_{i=1}^m w_i^2 \|y'(t_i) - f(y(t_i), t_i, s)\|^2, \quad (2.3)$$

where w_i^2 is the integration weight associated with time t_i ; we write the weights as squared quantities to avoid the cumbersome square root signs later on. Note that the time discretization of (2.3) may be a subset of the time grid used to integrate $x(t)$. In fact, the t_i may be chosen to select a small time window within $[0, T]$. Again, it is clear that if $\phi_m(y) = 0$, then $y(t_i) = x(t_i)$.

In what follows, we describe a method for approximating the process $x(t) = x(t, s)$ using a set of solutions computed at other points in the input parameter space. The essential idea is to construct an affine combination of these precomputed solutions, where the coefficients are computed with a nonlinear least squares minimization procedure on the discretized residual $\phi_m(y)$. It may not be immediately obvious that the approximation interpolates the precomputed solutions; we will justify the label *interpolant* with Theorem 4.2.

Let $x_j(t) = x(t, s_j)$ be the time dependent process with input parameters $s_j \in \mathcal{S}$ for $j = 1, \dots, n$; these represent the precomputed evaluations of (2.1) for the input parameters s_j , and we refer to them as the *bases*. To approximate the process $x(t) = x(t, s)$ for a given s , we seek constants $a_j = a_j(s)$ that are independent of time such that

$$x(t) \approx \tilde{x}(t) = \sum_{j=1}^n a_j x_j(t) \equiv X(t)a, \quad (2.4)$$

where a is an n -vector whose j th element is a_j , and $X(t)$ is a time dependent matrix whose j th column is $x_j(t)$. By definition, the coefficients of the affine combination satisfy

$$1 = \sum_{j=1}^n a_j = e^T a, \quad (2.5)$$

where e is an n -vector of ones. This constraint ensures that the approximation exactly reproduces components of $x(t)$ that do not depend on the parameters s . To see this,

let $\chi(t)$ be some component of the state vector $x(t, s)$ that is independent of s . Then

$$\sum_{j=1}^n a_j \chi(t) = \chi(t) \left(\sum_{j=1}^n a_j \right) = \chi(t). \quad (2.6)$$

For example, (2.5) guarantees that the approximation satisfies parameter independent boundary conditions, which arise in many applications of interest. Since the coefficients a are independent of time, the time derivative of the approximation can be computed as

$$\begin{aligned} \tilde{x}' &= \sum_{j=1}^n a_j x'_j \\ &= \sum_{j=1}^n a_j f(x_j(t), t, s_j) \\ &= F(t)a, \end{aligned}$$

where $F(t)$ is a time dependent matrix whose j th column is $f(x_j, t, s_j)$. Note that this can be modified to include a constant or time dependent mass matrix, as well.

With an eye toward computation, define the matrices $X_i = X(t_i)$ and $F_i = F(t_i)$. Then the discretized residual $\phi_m(\tilde{x})$ becomes

$$\begin{aligned} \phi_m(\tilde{x}) &= \sum_{i=1}^m w_i^2 \|\tilde{x}'(t_i) - f(\tilde{x}(t_i), t_i, s)\|^2 \\ &= \sum_{i=1}^m w_i^2 \|F_i a - f(X_i a, t_i, s)\|^2 \\ &\equiv \rho(a) \end{aligned}$$

To compute the coefficients a of the approximation, we solve the nonlinear least squares problem

$$\begin{aligned} &\underset{a}{\text{minimize}} && \rho(a) \\ &\text{subject to} && e^T a = 1. \end{aligned} \quad (2.7)$$

Note that the minimizer of (2.7) may not be unique due to potential nonconvexity in f , but this should not deter us. Since we are interested in approximating $x(t)$, any minimizer – or near minimizer – of (2.7) will be useful.

One could use a standard optimization routine [20] for a nonlinear objective with linear equality constraints to solve (2.7). However, we offer some specifics of a nonlinear least squares algorithm tuned to the details of this particular problem, namely (i) a single linear equality constraint representing the sum of the vector elements, (ii) the use of only evaluations of the forcing function to construct the data of the minimization problem, and (iii) the ill-conditioned nature of the problem.

2.1. A Nonlinear Least Squares Solver. To simplify the notation, we define the following quantities:

$$F = \begin{bmatrix} w_1 F_1 \\ \vdots \\ w_m F_m \end{bmatrix} \quad \varphi(a) = \begin{bmatrix} w_1 f(X_1 a, t_1, s) \\ \vdots \\ w_m f(X_m a, t_m, s) \end{bmatrix} \quad (2.8)$$

The residual $h : \mathbb{R}^n \rightarrow \mathbb{R}^{mp}$ is defined as

$$h(a) = Fa - \varphi(a) \quad (2.9)$$

so that the objective function $\rho(a)$ from (2.7) can be written as

$$\rho(a) = \|h(a)\|^2. \quad (2.10)$$

Let $\mathcal{J} = \mathcal{J}(a) \in \mathbb{R}^{mp \times n}$ be the Jacobian of $h(a)$. Given a guess $a_k \in \mathbb{R}^n$ such that $e^T a_k = 1$, the standard Newton step is computed by solving the constrained least squares problem

$$\begin{aligned} & \underset{\delta}{\text{minimize}} && \|\mathcal{J}_k \delta + h(a_k)\| \\ & \text{subject to} && e^T \delta = 0, \end{aligned} \quad (2.11)$$

where $\mathcal{J}_k = \mathcal{J}(a_k)$. Let δ_k be the minimizer, so that the update becomes

$$a_{k+1} = a_k + \delta_k. \quad (2.12)$$

The constraint on δ ensures that a_{k+1} sums to one.

Instead of using the standard Newton step (2.11), we wish to rewrite the problem slightly. Writing it in this alternative form suggests a method for dealing with the ill-conditioned nature of the problem, which we will explore in Section 5.2. We plug the update step (2.12) directly into the residual vector and exploit the constraint $e^T a_{k+1} = 1$ as

$$\begin{aligned} \mathcal{J}_k \delta_k + h(a_k) &= \mathcal{J}_k(a_{k+1} - a_k) + h(a_k) \\ &= \mathcal{J}_k a_{k+1} + (h(a_k) - \mathcal{J}_k a_k) \\ &= \mathcal{J}_k a_{k+1} + (h(a_k) - \mathcal{J}_k a_k) e^T a_{k+1} \\ &\equiv R_k a_{k+1} \end{aligned}$$

where

$$R_k = \mathcal{J}_k + (h(a_k) - \mathcal{J}_k a_k) e^T. \quad (2.13)$$

Written in this way, each Newton iterate a_{k+1} can be computed by solving the constrained least squares problem

$$\begin{aligned} & \underset{a}{\text{minimize}} && \|R_k a\| \\ & \text{subject to} && e^T a = 1. \end{aligned} \quad (2.14)$$

This is the heart of the residual minimizing model interpolation scheme. Notice that we used the standard Newton step in (2.12) without any sort of globalizing step length [9]. To include such a globalizer, we simply solve 2.14 for a_{k+1} and compute δ_k from (2.12).

We can relate the minimum residual of (2.14) to the true equation residual. Define $r_{k+1} = R_k a_{k+1}$, then

$$\|r_{k+1}\| \leq \|\mathcal{J}_k\| \|a_{k+1} - a_k\| + \sqrt{\rho(a_k)}. \quad (2.15)$$

Near the solution, we expect \mathcal{J}_k to be bounded and the difference in iteration to be small, so that the first term is negligible. We will use this bound to relate the norm of the equation residual to the conditioning of the constrained least squares problem in Theorem 4.6.

2.1.1. Jacobian-free Newton Step. To enable rapid implementation, we next show how to construct a finite difference Jacobian for the nonlinear least squares using only evaluations of the forcing function f from (2.1). We can write out \mathcal{J} as

$$\mathcal{J} = \nabla_a h = F - \mathcal{J}_f X, \quad (2.16)$$

where

$$\mathcal{J}_f = \begin{bmatrix} w_1 \nabla_x f(X_1 a, t_1, s) & & \\ & \ddots & \\ & & w_m \nabla_x f(X_m a, t_m, s) \end{bmatrix} \quad X = \begin{bmatrix} X_1 \\ \vdots \\ X_m \end{bmatrix}. \quad (2.17)$$

Notice that the Jacobian of h with respect to a contains terms with the Jacobian of f with respect to the states x multiplied by the basis vectors. Thus, we need only the action of \mathcal{J}_f on vectors, similar to Jacobian-free Newton-Krylov methods [14]. We can approximate the action of the Jacobian of f on a vector with a finite difference gradient; let x_j be the j th column of X , then

$$\mathcal{J}_f x_j \approx \frac{1}{\varepsilon} \left(\begin{bmatrix} w_1 f(X_1 a + \varepsilon x_j(t_1), t_1, s) \\ \vdots \\ w_m f(X_m a + \varepsilon x_j(t_m), t_m, s) \end{bmatrix} - \begin{bmatrix} w_1 f(X_1 a, t_1, s) \\ \vdots \\ w_m f(X_m a, t_m, s) \end{bmatrix} \right). \quad (2.18)$$

We can assume that the terms $f(X_i a, t_i, s)$ were computed before approximating the Jacobian to check the norm of the residual. Therefore, at each Newton iteration we need nm evaluations of f to compute the approximate Jacobian – one for each basis at each point in the time discretization. The implementation will use this finite difference approximation. But for the remainder of the analysis, we assume we have the true Jacobian.

2.1.2. Initial Guess. The convergence of nonlinear least squares methods depends strongly on the initial guess. In this section, we propose an initial guess based on treating the forcing function f as though it was linear in the states. To justify this treatment, let $f(x) = f(x, t, s)$ for given t and s . We take the Taylor expansion about the approximation \tilde{x} as

$$f(x) = f(\tilde{x}) + \nabla_x f(\tilde{x})(x - \tilde{x}) + \dots \quad (2.19)$$

Evaluate this expansion at $x_j = x_j(t)$, multiply it by a_j , and sum over j to get

$$\begin{aligned} \sum_{j=1}^n a_j f(x_j) &= \sum_{j=1}^n a_j f(\tilde{x}) + \sum_{j=1}^n a_j \nabla_x f(\tilde{x})(x_j - \tilde{x}) + \dots \\ &= f(\tilde{x}) \underbrace{\left(\sum_{j=1}^n a_j \right)}_{=1} + \nabla_x f(\tilde{x}) \underbrace{\left(\sum_{j=1}^n a_j x_j - \tilde{x} \right)}_{=0} + \dots \end{aligned}$$

The constraint eliminates the first order terms in Taylor expansion. If we ignore the higher order terms, then we can approximate

$$f(\tilde{x}, t, s) \approx \sum_{j=1}^n a_j f(x_j, t, s) \equiv G(t)a, \quad (2.20)$$

where the j th column of $G(t)$ is $f(x_j, t, s)$. Let $G_i = G(t_i)$, and define the $mp \times n$ matrix G as

$$G = \begin{bmatrix} w_1 G_1 \\ \vdots \\ w_m G_m \end{bmatrix}. \quad (2.21)$$

Then to solve for the initial guess a_0 , we define

$$R_{-1} = F - G \quad (2.22)$$

and solve (2.14). In words, we treat f as linear in the states and solve the same constrained least squares problem. Of course, if f is actually linear, then this is the only step in the approximation; no Newton iterations are required. In fact, we show in Theorem 4.3 that using this procedure alone to compute the coefficients a also produces an interpolant. Therefore, if the system is locally close to linear in the states, then a small number of Newton iterations will be sufficient.

3. Comparison to Kriging Interpolation. In this section, we compare the quantities computed in the residual minimizing scheme to kriging interpolation, which is commonly used in geostatistics. Suppose that $x(t, s)$ is a scalar valued process with $t \in [0, T]$ and $s \in \mathcal{S}$. Following kriging nomenclature, we treat $[0, T]$ as the sample space and t as a random coordinate. Assume for convenience that $x(t, s)$ has mean zero at each $s \in \mathcal{S}$,

$$\mathbb{E}[x(s)] = \int_0^T x(t, s) dt = 0, \quad s \in \mathcal{S}. \quad (3.1)$$

The covariance function of the process is then

$$\text{cov}(s_i, s_j) = \mathbb{E}[x(s_i)x(s_j)]. \quad (3.2)$$

Next suppose we are given the fixed values $x_j = x(s_j)$, and we wish to approximate $x = x(s)$ with the affine model

$$x \approx \tilde{x} = \sum_{j=1}^n x_j a_j, \quad \sum_{j=1}^n a_j = 1. \quad (3.3)$$

This is the model used in ordinary kriging interpolation [8]. To compute the coefficients a_j , one builds the following linear system of equations from the assumed covariance function (3.2) (which is typically a model calibrated with the $\{s_j, x_j\}$ pairs). Define $A_{ij} = \text{cov}(s_i, s_j)$ and $b_i = \text{cov}(s_i, s)$. Then the vector of coefficients a satisfies

$$\begin{bmatrix} A & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} a \\ \lambda \end{bmatrix} = \begin{bmatrix} b \\ 1 \end{bmatrix}, \quad (3.4)$$

where λ is the Lagrange multiplier.

Next we compare the kriging approach to model interpolation on a scalar valued process. Let x_j be the vector whose i th element is $x(t_i, s_j)$ with $i = 1, \dots, m$; in other words, x_j contains the time history of the process x with input parameters s_j . Again,

assume that the temporal average of $x(t, s)$ is zero for all $s \in \mathcal{S}$. The affine model to approximate the time history for $x = x(s)$ is

$$x \approx \tilde{x} = \sum_{j=1}^n x_j a_j = Xa, \quad \sum_{j=1}^n a_j = e^T a = 1, \quad (3.5)$$

where X is a matrix whose j th column is x_j . Notice the similarity between (3.3) and (3.5). One way to compute the coefficients a in the spirit of an error minimizing scheme would be to solve the least squares problem

$$\begin{aligned} & \underset{a}{\text{minimize}} && \|Xa - x\| \\ & \text{subject to} && e^T a = 1. \end{aligned} \quad (3.6)$$

Ignore for the moment that solving this least squares problem requires x – the exact vector we are trying to approximate. The KKT system associated with the constrained least squares problem is

$$\begin{bmatrix} X^T X & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} a \\ \lambda \end{bmatrix} = \begin{bmatrix} X^T x \\ 1 \end{bmatrix}. \quad (3.7)$$

But notice that each element of $X^T X$ is an inner product between two time histories, and this can be interpreted as approximating the covariance function with an empirical covariance, i.e.

$$x_k^T x_j = \sum_{i=1}^m x(t_i, s_k) x(t_i, s_j) \approx (m-1) \text{cov}(s_k, s_j). \quad (3.8)$$

Similarly for the point s ,

$$x_k^T x = \sum_{i=1}^m x(t_i, s_k) x(t_i, s) \approx (m-1) \text{cov}(s_k, s). \quad (3.9)$$

Then,

$$\frac{1}{m-1} X^T X \approx A, \quad \frac{1}{m-1} X^T x \approx b, \quad (3.10)$$

where A and b are from (3.4). We can multiply the top equations of (3.7) by $1/(m-1)$ to transform it to an empirical form of (3.4). Loosely speaking, the residual minimizing method for computing the coefficients a is a transformed version of the least squares problem (3.6), where the transformation comes from the underlying model equations. In this sense, we can think of the residual minimizing method as a version of kriging where the covariance information comes from sampling the underlying dynamical model.

4. Analysis. We begin this section with a summary its results. We first examine the quality of a linear approximation of the process and derive a lower bound for the average error over the input parameter space; such analysis is related to the error bounds found in POD-based model reduction [1, 22]. We then show that the residual minimizing model interpolates the problem data; in fact, even the initial guess proposed in Section 2.1.2 has an interpolatory property in most cases. We then show that the coefficients inherit the type of input parameter dependence from the underlying

dynamical model. In particular, if the dynamical system forcing function depends continuously on the input parameters, then so do the coefficients; the interpolatory property and continuity of the coefficients are appealing aspects of the scheme. Next we prove that each added basis improves the approximation over the entire parameter space. Finally, the last theorem of the section relates the minimum singular value of the data matrix in the Newton step to the equation residual; this result implies that to achieve an approximation with a small residual one must solve an ill-conditioned least squares problem.

4.1. An Error Bound. Before addressing the characteristics of the interpolant, we may ask how well a linear combination of basis vectors can approximate some parameterized vector. The following theorem gives a lower bound on the average error in the best n -term linear approximation in the Euclidean norm. This bound is valid for any linear approximation of a parameterized vector – not necessarily the time history of a parameterized dynamical system – and thus applies to any linear method, including those mentioned in the introduction. The lower bound is useful for determining a best approximation. If our approximation behaves like the lower bound as n increases, then we can claim that it behaves like the best approximation.

THEOREM 4.1. *Let $x = x(s)$ be a parameterized p -vector, and let X be a full rank matrix of size $p \times n$ with $n < p$. Define the symmetric, positive semi-definite matrix*

$$C = \int_S x x^T ds \quad (4.1)$$

and let $C = U\Theta U^T$ be its eigenvalue decomposition, where $\theta_1 \geq \dots \geq \theta_p$ are the ordered eigenvalues. Then

$$\int_S \min_a \|Xa - x\|^2 ds \geq \sqrt{\sum_{k=n+1}^p \theta_k^2}. \quad (4.2)$$

Proof. For a fixed s with $x = x(s)$, we examine the least squares problem

$$\min_a \|Xa - x\|. \quad (4.3)$$

Using the normal equations, we have

$$a = (X^T X)^{-1} X^T x. \quad (4.4)$$

Then the minimum residual is given by

$$\min_a \|Xa - x\| = \|(X(X^T X)^{-1} X^T - I)x\| = \|Bx\|, \quad (4.5)$$

where $B = X(X^T X)^{-1} X^T - I$. Denote the Frobenius norm by $\|\cdot\|_F$. Taking the average of the minimum norm squared, we have

$$\int_S \|Bx\|^2 ds = \int_S x^T B^T B x ds \quad (4.6)$$

$$= \int_S \|Bx x^T B^T\|_F ds \quad (4.7)$$

$$\geq \left\| B \left(\int_S x x^T ds \right) B^T \right\|_F \quad (4.8)$$

$$= \|BCB^T\|_F, \quad (4.9)$$

where (4.8) comes from Jensen's inequality. Suppose that X contains the first n columns of U , i.e. partition

$$U = \begin{bmatrix} X & Y \end{bmatrix}. \quad (4.10)$$

Similarly partition the associated eigenvalues

$$\Theta = \begin{bmatrix} \Theta_1 & \\ & \Theta_2 \end{bmatrix}. \quad (4.11)$$

Then since X and Y are orthogonal,

$$\begin{aligned} B &= X(X^T X)^{-1} X^T - I \\ &= X X^T - I \\ &= -Y Y^T. \end{aligned}$$

In this case

$$\begin{aligned} \|BCB^T\|_F &= \|Y Y^T C Y Y^T\|_F \\ &= \|Y \Theta_2 Y^T\|_F \\ &= \sqrt{\sum_{k=n+1}^p \theta_k^2} \end{aligned}$$

Therefore, for a given X (not necessarily the eigenvectors),

$$\int_S \min_a \|Xa - x\|^2 ds \geq \|BCB^T\|_F \geq \sqrt{\sum_{k=p+1}^n \theta_k^2}, \quad (4.12)$$

as required. \square

In words, Theorem 4.1 states that the average optimal linear approximation error is at least as large as the norm of the neglected eigenvalues of covariance-like matrix C . As an aside, we note that if $p = n$, then the best approximation error is zero, since X is an invertible matrix.

The result in Theorem 4.1 is similar in spirit to the approximation properties of POD-Galerkin based model reduction techniques [1] and the best approximation results for Karhunen-Loeve type decompositions [17]. Given a matrix of snapshots X whose j th column is $x(s_j)$, one can approximate the matrix $C \approx n^{-1} X X^T$; the error bounds for POD-based reduced order models are typically given in terms of the singular values of X . We can approximate this lower bound for a given problem and compare it to the error for the residual minimizing reduced model; we will see one numerical example that the approximation error behaves like the lower bound.

4.2. Interpolation and Continuity. A few important properties are immediate from the construction of the residual minimizing model. Existence follows from existence of a minimizer for the nonlinear least squares problem (2.7). Also, by construction, the coefficients a provide the optimal approximation in the space spanned by the bases, where optimality is with respect to the surrogate error measure given by the objective function of (2.7) – i.e., the residual – under the constraint that $e^T a = 1$. As in other residual minimizing schemes, the minimum value of the objective function

provides an a posteriori error measure for the approximation. The next few theorems expose additional interesting properties of the residual minimizing approximation.

THEOREM 4.2. *The residual minimizing approximation $\tilde{x}(t, s)$ interpolates the basis elements, i.e.,*

$$\tilde{x}(t_i, s_j) = x_j(t_i), \quad i = 1, \dots, m, \quad j = 1, \dots, n. \quad (4.13)$$

Proof. Let a be a vector such that $e^T a = 1$ and

$$X_i a = x_j(t_i), \quad F_i a = f(x_j(t_i), t_i, s_j), \quad i = 1, \dots, m. \quad (4.14)$$

Such a vector exists; the n -vector of zeros with 1 in the j th entry satisfies these conditions. For $\rho(a) = \rho(a, s_j)$ from (2.7),

$$\begin{aligned} \rho(a) &= \sum_{i=1}^m w_i^2 \|F_i a - f(X_i a, t_i, s_j)\|^2 \\ &= \sum_{i=1}^m w_i^2 \|f(x_j(t_i), t_i, s_j) - f(x_j(t_i), t_i, s_j)\|^2 = 0 \end{aligned}$$

which is a minimum. Then

$$\tilde{x}(t_i, s_j) = X_i a = x_j(t_i), \quad i = 1, \dots, m. \quad (4.15)$$

Since j was arbitrary, this completes the proof. \square

Theorem 4.2 justifies the label *interpolant* for the approximation $\tilde{x}(t, s)$. For many cases, this property extends to the initial guess described in Section 2.1.2.

THEOREM 4.3. *Let $R_{-1} = R_{-1}(s)$ be defined as in (2.22). Assume the matrix $[R_{-1}^T, e]^T$ has full column rank for all $s = s_j$. Then the initial guess, denoted by $\tilde{x}_0(t, s)$, interpolates the basis elements, i.e.,*

$$\tilde{x}_0(t_i, s_j) = x_j(t_i), \quad i = 1, \dots, m, \quad j = 1, \dots, n. \quad (4.16)$$

Proof. The full rank assumption on $[R_{-1}^T, e]^T$ at each $s = s_j$ implies that the constrained least squares problem

$$\begin{aligned} &\underset{a}{\text{minimize}} && \|R_{-1} a\| \\ &\text{subject to} && e^T a = 1 \end{aligned} \quad (4.17)$$

has a unique solution. Let a be a vector of zeros with 1 in the j th element. Then $e^T a = 1$ and

$$\begin{aligned} \|R_{-1} a\| &= \|F a - G a\| \\ &= \left\| \begin{bmatrix} w_1 f(x_j(t_1), t_1, s_j) \\ \vdots \\ w_m f(x_j(t_m), t_m, s_j) \end{bmatrix} - \begin{bmatrix} w_1 f(x_j(t_1), t_1, s_j) \\ \vdots \\ w_m f(x_j(t_m), t_m, s_j) \end{bmatrix} \right\| = 0, \end{aligned}$$

which is a minimum. By the uniqueness,

$$\tilde{x}_0(t_i, s_j) = X_i a = x_j(t_i), \quad i = 1, \dots, m. \quad (4.18)$$

Since j was arbitrary, this completes the proof. \square

The proof required the full rank assumption on $[R_{-1}^T, e]^T$ at each $s = s_j$. We expect this to be the case in practice with a properly selected bases x_j . If we extend this assumption to all $s \in \mathcal{S}$, then we make a statement about the how the coefficients of the interpolant behave over the parameter space. In the next theorem, we show that the coefficients of the residual minimizing interpolant inherit the behavior of the model terms with respect to the input parameters.

THEOREM 4.4. *Let $R_k = R_k(s)$ be defined as in (2.13), and assume the matrix $[R_k^T, e]^T$ has full column rank for all $s \in \mathcal{S}$. If the function $f(x, t, s)$ and its Jacobian $\nabla_x f$ depend continuously on s , then the coefficients $a(s)$ computed with a finite number of Newton iterations are also continuous with respect to s .*

Proof. We can write the KKT system associated with the constrained least squares problem (2.14) as

$$\begin{bmatrix} R_k^T R_k & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} a \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}, \quad (4.19)$$

where $\lambda = \lambda(s)$ is the Lagrange multiplier. The full rank assumption on $[R_k^T, e]^T$ implies that the KKT matrix is invertible for all $s \in \mathcal{S}$, and we can write

$$\begin{bmatrix} a \\ \lambda \end{bmatrix} = \begin{bmatrix} R_k^T R_k & e \\ e^T & 0 \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}. \quad (4.20)$$

Since the elements of R_k are continuous in s , so are the elements of the inverse of the KKT system, which implies that $a(s)$ is continuous in s , as required. \square

4.3. Convergence and Conditioning. The convergence analysis we present is somewhat nonstandard. As opposed to computing an a priori rate of convergence, we show in the next theorem that the minimum residual decreases monotonically with each added basis.

THEOREM 4.5. *Let*

$$X^{(n)} = \begin{bmatrix} X_1 \\ \vdots \\ X_m \end{bmatrix} \quad (4.21)$$

be a given basis with n columns, and let a_n^ be the minimizer of (2.14) with associated minimum function value $\rho_n^* = \rho(a_n^*, s)$. Let $s_{n+1} \in \mathcal{S}$ and*

$$x_{n+1} = \begin{bmatrix} x(t_1, s_{n+1}) \\ \vdots \\ x(t_m, s_{n+1}) \end{bmatrix}. \quad (4.22)$$

Define the updated basis

$$X^{(n+1)} = \begin{bmatrix} X^{(n)} & x_{n+1} \end{bmatrix}. \quad (4.23)$$

Let a_{n+1}^ be the minimizer of (2.14) with the basis $X^{(n+1)}$, and let $\rho_{n+1}^* = \rho(a_{n+1}^*, s)$ be the minimum function value. Then there is an $0 \leq \alpha \leq 1$ such that*

$$\rho_{n+1}^* = \alpha \rho_n^*, \quad (4.24)$$

for all $s \in \mathcal{S}$.

Proof. Define

$$a = \begin{bmatrix} a_n^* \\ 0 \end{bmatrix}, \quad (4.25)$$

and note that $e^T a = 1$. Then

$$\rho_{n+1}^* \leq \rho(a, s) = \rho_n^*. \quad (4.26)$$

Next let $s = s_{n+1}$, and let a be a $n+1$ -vector of zeros with a 1 in the last entry. Then

$$\rho_{n+1}^* = \rho(a, s_{n+1}) = 0, \quad (4.27)$$

as required. \square

In words, the minimum residual decreases monotonically as bases are added to the approximation. Since s is arbitrary, Theorem 4.5 says that each basis added to the approximation reduces the residual norm *globally* over the parameter space; in the worst case, it does no harm, and in the best case it achieves the true solution. This result is similar to Theorem 2.2 from [5].

In the final theorem of this section, we show that achieving a small residual requires the solution of an ill-conditioned constrained least squares problem for the Newton step. To do this, we first reshape the constrained least squares problem using the standard null space method [4], and then we apply a result from [25] relating the minimum norm of the residual to the minimum singular value of the data matrix.

THEOREM 4.6. *Let R_k be the matrix from the constrained least squares problem for the Newton step (2.14), and define $\sigma_{\min}(R_k)$ to be its minimum singular value. Then*

$$\sigma_{\min}(R_k) \leq \sqrt{n} \left(\|\mathcal{J}\| \|a_{k+1} - a_k\| + \sqrt{\rho(a_k)} \right). \quad (4.28)$$

Proof. Let $R = R_k$ be defined as in (2.13). To solve the constrained least squares problem (2.14) via the nullspace method, we take a QR factorization of the constraint vector

$$Q^T e = \sqrt{n} e_1, \quad (4.29)$$

where e_1 is an n -vector of zeros with a one in the first entry. Partition $Q = [q_1, Q_2]$, where q_1 is the first column of Q . The $n \times (n-1)$ matrix Q_2 is a basis for the null space of the constraint. It can be shown that $q_1 = n^{-1/2} e$. Using this transformation, the constrained least squares problem becomes the unconstrained least squares problem

$$\underset{v}{\text{minimize}} \quad \|RQ_2 v + \frac{1}{n} R e\|. \quad (4.30)$$

Then the Newton update is given by

$$a_{k+1} = \frac{1}{n} e + Q_2 v^*, \quad (4.31)$$

where v^* is the minimizer of (4.30). For a general overdetermined least squares problem $\min_u \|Au - b\|$, Van Huffel [25] shows that for $\gamma > 0$,

$$\frac{\sigma_{\min}([b\gamma, A])}{\gamma} \leq \|b - Au^*\|, \quad (4.32)$$

where u^* is the minimizer; we apply this result to (4.30). Taking $\gamma = \sqrt{n}$, we have

$$\left[\frac{1}{\sqrt{n}} Re, RQ_2 \right] = RQ. \quad (4.33)$$

Since Q is orthogonal, $\sigma_{\min}(RQ) = \sigma_{\min}(R)$. Then

$$\sigma_{\min}(R) \leq \sqrt{n} \left\| RQ_2 v^* + \frac{1}{n} Re \right\| = \sqrt{n} \|Ra_{k+1}\|. \quad (4.34)$$

Combining this result with the bound on the residual (2.15) achieves the desired result. \square

To reiterate, near the solution of the nonlinear least squares problem (2.7), we expect $\|\mathcal{J}\|$ to be bounded and the difference $\|a_{k+1} - a_k\|$ to be small so that the first term in the bound becomes negligible. Therefore, a solution with a small residual norm $\rho(a)$ will require the solution of a constrained least squares problem with a small minimum singular value, which implies a large condition number. To combat this, we propose some heuristics for dealing with the ill-conditioning of the problem in the following section.

5. Computational Heuristics. In this section, we propose heuristics for (1) reducing the cost of constructing the interpolant, (2) mitigating the ill-conditioning of the least squares problems, and (3) adding new bases to the approximation.

5.1. Cost Reduction. The reader may have noticed that, if we measure the cost of the residual minimizing scheme in terms of number of evaluations of f from (2.1) (assuming we use the finite difference Jacobian described in Section 2.1.1), then computing the approximation can be more expensive than evaluating the true model. Specifically, an explicit one-step time stepping scheme evaluated at the time discretization t_1, \dots, t_m requires m evaluations of f . However, merely constructing the matrix G in (2.22) at the same time discretization to compute the initial guess requires mn evaluations of f for n bases. And if the elements of the matrix F were not stored during the runs, computing them requires another mn evaluations of f . On top of that, each Newton step requires yet another mn evaluations. This simple cost analysis would seem to discourage us from comparing the interpolation method to other methods for model reduction.

The interpolation method becomes an appropriate method for model reduction when one is interested in a small subset of the time domain. For example, if a quantity of interest is computed as a function of the state at some final time, then the interpolation method can be used to approximate the state at the final time at a new parameter value without computing the full history. This is particularly useful for dynamical systems with rapidly varying initial transients that require small time steps. With the interpolation method, one need not resolve the initial transients with a small time step to approximate the state at the final time. Instead, the method takes advantage of regularity in the parameter space to construct the interpolating approximation.

Alternatively, if one were interested in a minimal set of points in time for coarse approximation of the history, these could be determined with a method similar to the discrete empirical interpolation method [7] in the time domain. With this set of discretization points, one could approximate the time history at a new parameter value without the full model solver. We do not pursue this idea further in this work, but we believe it holds promise for approximating time-averaged quantities of interest.

This sort of reduction is applicable to the number of points m in the time domain. In the parameter domain, the number of points n corresponds to the number of full model solutions used in the affine combination (2.4). Typically, we think that n is small due to the cost of the full model solution. However, in cases where a state vector may be represented by an even smaller set of basis vectors, we can borrow an idea from moving (or windowed) least squares [27] to reduce the number of columns in the solves for the Newton steps (2.14).

Choose an integer $M < n$. For a parameter point s , find the M points in the set of $\{s_j\}$ that are nearest to s . Then use only the x_j corresponding to nearby s_j to construct the least squares approximations. In the numerical examples, we demonstrate the savings generated by this heuristic.

5.2. Alleviating Ill-Conditioning. In Theorem 4.6, we showed that achieving a small residual required the solution of an ill-conditioned least squares problem for the Newton step. Here we offer a heuristic for alleviating the effects of the ill-conditioning. The heart of the residual minimizing scheme is the constrained linear least squares problem (2.14) which we rewrite with a general $m \times n$ matrix R as

$$\begin{aligned} & \underset{a}{\text{minimize}} && \|Ra\| \\ & \text{subject to} && e^T a = 1. \end{aligned} \quad (5.1)$$

For now we assume that R is full rank, although it may have a very large condition number. The particular form of the least squares problem (the single linear constraint e and the zero right hand side) will permit us to use some novel approaches for dealing with the ill-conditioning.

To analyze this problem, we first derive a few useful expressions. Let λ be the Lagrange multiplier associated with the constraint. The minimizer $[a^T, -\lambda]^T$ satisfies the KKT conditions

$$\begin{bmatrix} R^T R & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} a \\ -\lambda \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (5.2)$$

We can use a Schur complement (or block elimination) method to derive expressions for the solution:

$$\lambda = \frac{1}{e^T (R^T R)^{-1} e}, \quad a = \lambda (R^T R)^{-1} e. \quad (5.3)$$

Also note that the first KKT equation gives

$$R^T R a = \lambda e. \quad (5.4)$$

Premultiplying this by a^T , we get

$$a^T R^T R a = \lambda \underbrace{a^T e}_{=1}, \quad (5.5)$$

or equivalently

$$\|Ra\|^2 = \lambda. \quad (5.6)$$

In other words, the value of the optimal Lagrange multiplier is the squared norm of the residual. We can use this fact to improve the conditioning of computing the approximation while maintaining a small residual.

Since we wish to minimize $\|Ra\|$, it is natural to look for a linear combination of the right singular vectors of R associated with the smallest singular values. We first compute the thin singular value decomposition (SVD) of R as

$$R = U\Sigma V^T \quad \text{diag}(\Sigma) = [\sigma_1, \dots, \sigma_n], \quad (5.7)$$

and rotate (5.1) by the right singular vectors. Define

$$\hat{a} = V^T a, \quad d = V^T e. \quad (5.8)$$

Then (5.1) becomes

$$\begin{aligned} & \underset{a}{\text{minimize}} && \|\Sigma \hat{a}\| \\ & \text{subject to} && d^T \hat{a} = 1, \end{aligned} \quad (5.9)$$

with associated KKT conditions

$$\begin{bmatrix} \Sigma^2 & d \\ d^T & 0 \end{bmatrix} \begin{bmatrix} \hat{a} \\ -\lambda \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (5.10)$$

Applying block elimination to solve this system involves inverting Σ^2 and multiplying by d – an operation with a condition number of $(\sigma_1/\sigma_n)^2$. We can improve the conditioning by solving a truncated version of the problem. Let $k < n$ be a truncation and partition

$$\Sigma = \begin{bmatrix} \Sigma_1 & \\ & \Sigma_2 \end{bmatrix} \quad V = [V_1 \quad V_2] \quad d = \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} \quad \hat{a} = \begin{bmatrix} \hat{a}_1 \\ \hat{a}_2 \end{bmatrix} \quad (5.11)$$

according to k . Since we want to minimize the residual $\|\Sigma \hat{a}\|$ and the largest values of Σ appear at the top, we set $\hat{a}_1 = 0$ and solve

$$\begin{bmatrix} \Sigma_2^2 & d_2 \\ d_2^T & 0 \end{bmatrix} \begin{bmatrix} \hat{a}_2^{(t)} \\ -\lambda^{(t)} \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad (5.12)$$

where the superscript (t) is for *truncated*; the condition number of inverting and multiplying by Σ_2^2 is $(\sigma_{k+1}/\sigma_n)^2$. Then the linear combination of the right singular vectors associated with the smallest singular values is

$$a^{(t)} = V_2^T \hat{a}_2^{(t)}. \quad (5.13)$$

We can measure what was lost in the truncation in terms of the minimum residual by looking at the difference between the Lagrange multipliers

$$\left| \|Ra\|^2 - \|Ra^{(t)}\|^2 \right| = |\lambda - \lambda^{(t)}|. \quad (5.14)$$

And this suggests a way to choose the truncation a priori. Once the SVD of R has been computed, we can trivially check the minimum residual norm by computing the Lagrange multipliers for various truncations; see (5.6). We want to truncate to reduce the condition number of the problem and avoid disastrous numerical errors. But we do not want to truncate so much that we lose accuracy in the approximation. Therefore, we propose the following. For $j = 1, \dots, n$, compute

$$\begin{aligned} d_j &= [d_{n-j+1}, \dots, d_n]^T && \text{the last } j \text{ elements of } d \\ \Sigma_j &= \text{diag}(\sigma_{n-j+1}, \dots, \sigma_n) && \text{the last } j \text{ singular values} \\ y_j &= \Sigma_j^{-1} d_j \\ \lambda_j &= \frac{1}{y_j^T y_j} \end{aligned}$$

The λ_j measure the norm of the residual for a truncation that retains the last j singular values; λ_n is the true minimum residual. Therefore, we seek the smallest k such that

$$|\lambda_k - \lambda_n| < \tau \quad (5.15)$$

for a given tolerance τ representing how much minimum residual we are willing to sacrifice for better conditioning, and this determines our truncation.

5.2.1. A Note on Rank Deficiency. There is an interesting tension in the constrained least squares problem (5.1). The ideal scenario would be to find a vector in the null space of R that also satisfies the constraint. This implies that we want R to be rank deficient for the sake of the approximation. However, if the matrix $[R^T, e]^T$ is rank deficient, then there are infinite solutions for the constrained least squares problem [4]. To distinguish between these two potential types of rank deficiency, we apply the following. Suppose R is rank deficient and we have partitioned

$$R = \begin{bmatrix} U_1 & U_2 \end{bmatrix} \begin{bmatrix} \Sigma_1 & \\ & 0 \end{bmatrix} \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix}. \quad (5.16)$$

Next compute $d = V_2^T e$. All components of d equal to zero correspond to singular vectors that are in the null space in the constraint; we want to avoid these. Suppose

$$d = \begin{bmatrix} d_1 \\ 0 \end{bmatrix} = \begin{bmatrix} V_{2,1}^T e \\ V_{2,2}^T e \end{bmatrix}. \quad (5.17)$$

Then we can choose any linear combination of the vectors $V_{2,1}$ that satisfies the constraint to achieve a zero residual; one simple choice is to take the component-wise average of the vectors $V_{2,1}$.

If all of $d = 0$, then we must apply a method for the rank deficient constrained least squares problem, such as the null space method to transform the (5.1) to an unconstrained problem coupled with the pseudoinverse to compute the minimum norm solution [4].

The drawback of this SVD-based approach is that each Newton step requires computing the SVD. And we expect that most parameter and uncertainty studies will require many such evaluations. Therefore we are actively pursuing more efficient methods for solving (5.1).

5.3. Adding Bases. When doing a parameter study on a complex engineering system, the question of where in the parameter space to compute the solution arises frequently. In the context of this residual minimizing model interpolation, we have a natural method and metric for answering this question. Let $\rho^* = \rho^*(s)$ be the minimum objective function value from the nonlinear least squares problem (2.7) with a given basis. Then the point $s_{n+1} \in \mathcal{S}$ to next evaluate the full model is the *maximizer* of

$$\underset{s \in \mathcal{S}}{\text{maximize}} \quad \rho^*(s), \quad (5.18)$$

so that $x_{n+1}(t) = x(t, s_{n+1})$. Of course, the optimization problem (5.18) is in general non-concave, and derivatives with respect to the parameters s are often not available. However, we are comforted by the fact that *any* point in the parameter space with a positive value for ρ^* will yield a basis element that will improve the approximation; see Theorem 4.5. We therefore expect derivative-free global optimization

heuristics [16] to perform sufficiently well for many applications. The idea of maximizing the residual over the parameter space to compute the next full model solution was also proposed in [5] in the context of reduced order modeling, and they present a thorough treatment of the resulting optimization problem – including related greedy sampling techniques [11].

There are many possible variations on (5.18). We have written the optimization to chose a single basis to add. However, multiple independent optimizations could be run, and a subset of the computed optima could be added to accelerate convergence.

6. Numerical Examples. We numerically study two models in this section: (i) a three-state nonlinear dynamical system representing a chemical kinetics mechanism with a single input parameter controlling the stiffness of the system, and (ii) a nonlinear, parabolic PDE modeling conductive heat transfer with a temperature dependent random field model of the thermal conductivity. The first problem is a toy model used to explore and confirm properties of the approximation scheme; we do not attempt any reduction in this case. The second example represents a step toward a large scale application in need of model reduction. All numerical experiments were performed on a dual quad-core Intel W5590 with 12GB of RAM running Ubuntu and MATLAB 2011b. Scripts for reproducing the experiments can be found at www.stanford.edu/~paulcon/rmmi.html; the second experiment requires the MATLAB PDE Toolbox.

6.1. A 3-Species Kinetics Problem. The following model from [24] contains the relevant features of a stiff chemical kinetics mechanism. Let

$$x(t, s) = \begin{bmatrix} u(t, s) \\ v(t, s) \\ w(t, s) \end{bmatrix} \quad (6.1)$$

be the three state variables with evolution described by the function

$$f(x, t, s) = \begin{bmatrix} -5\frac{u}{s} - \frac{uv}{s} + vw + 5\frac{v^2}{s^2} + \frac{w}{s} - u \\ 10\frac{u}{s} - \frac{uv}{s} - vw - 10\frac{v^2}{s^2} + \frac{w}{s} + u \\ \frac{uv}{s} - vw - \frac{w}{s} + u \end{bmatrix}, \quad (6.2)$$

where the parameter s controls the stiffness of the system; the initial state is $u_0 = v_0 = w_0 = 0.5$. We examine the parameter range $s \in [0.005, 1.2]$, where the smaller value of s corresponds to a stiffer system. Given s , we solve for the evolution of the states using MATLAB's `ode45` routine from $t = 0$ to $t = 1$. The ODE solver has its own adaptive time stepping method. We extract a discrete time process on 300 equally spaced points in the time interval $[0, 1]$; denote these points in time by t_j where

$$t_j = j\Delta t, \quad \Delta t = \frac{1}{300}. \quad (6.3)$$

In Figure 6.1, we plot the evolution of each component for $s = 0.005$ and $s = 1.2$. We can see that the stiffness of the equation yields rapidly varying initial transients for s near zero, which makes this a challenging problem.

We approximate the average error and average minimum residual over the parameter space by using 300 equally spaced points in the parameter range $[0.005, 1.2]$; denote these points by

$$s_k = 0.005 + k\Delta s, \quad \Delta s = \frac{1.195}{300}. \quad (6.4)$$

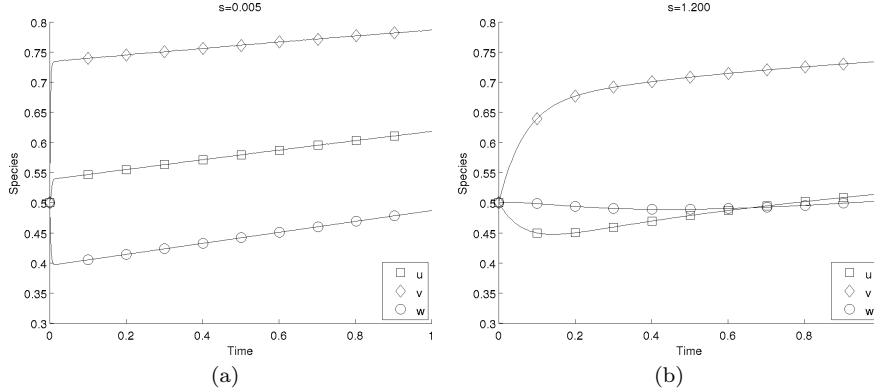


Fig. 6.1: The evolution of the states described by (6.2) for (a) $s = 0.005$ corresponding to a stiff system and (b) $s = 1.200$ corresponding to a nonstiff system.

The approximate average error and minimum residual are computed as

$$\mathcal{E} = \Delta s \Delta t \sum_k \sum_j \|\tilde{x}(t_j, s_k) - x(t_j, s_k)\|, \quad (6.5)$$

$$\mathcal{R} = \Delta s \Delta t \sum_k \sum_j \|\tilde{x}'(t_j, s_k) - f(\tilde{x}(t_j, s_k), s_k)\|. \quad (6.6)$$

We examine the reduction in \mathcal{E} and \mathcal{R} as more bases are added according to the heuristic in Section 5.3, and we compare that with the lower bound from Theorem 4.1. We compute the lower bound from Theorem 4.1 (the right hand side of (4.2)) using a Gauss-Legendre quadrature rule with 9,600 points in the range of s , which was sufficient to obtain converged eigenvalues.

We also use the discretization with s_k to approximate the optimization problem (5.18) used to select additional basis elements. In other words, for a given basis we compute the minimum residual at each s_k , and we append to the basis the true solution corresponding to the s_k that yields the largest minimum residual. We begin with 2 bases – one at each end point of the parameter space – and stop after 40 have been added. \mathcal{E} and \mathcal{R} appear in Figure 6.2 compared to the computed lower bound. We see that the error behaves roughly like the lower bound. For each evaluation of the reduced order model, we track the number of Newton iterations, and we plot the average over the 300 s_k in Figure 6.2. Notice that the number of necessary Newton iterations decreases as the number of basis functions increases.

To confirm the results of Theorem 4.6, we plot the minimum residual at each parameter point with the maximum condition number of the matrices R_k from (2.14) used to compute the Newton steps. In Figure 6.3 we show this plot for 5, 10, 20, and 40 bases. We clearly see an inverse relationship between the condition number and the minimum residual, which is consistent with Theorem 4.6.

6.2. Nonlinear Transient Heat Conduction Study. Next, we examine a two-dimensional transient heat conduction model for a steel beam with uncertain material properties in a high temperature environment. For a given realization of the thermal conductivity, the output of interest is the proportion of the domain whose

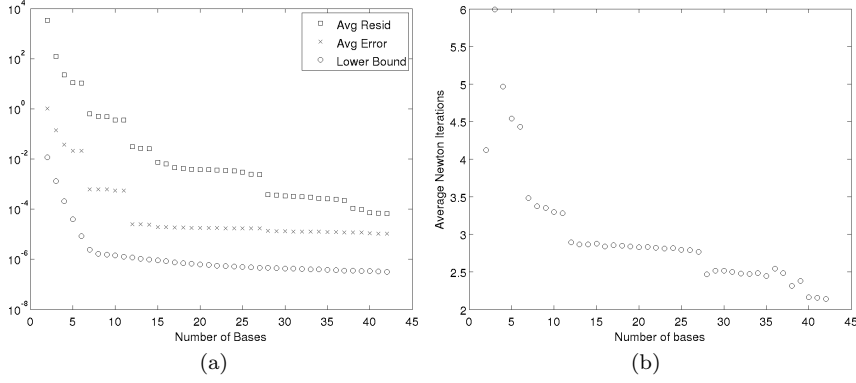


Fig. 6.2: (a) Reduction in average minimum residual \mathcal{R} (6.6), average error \mathcal{E} (6.5), and lower bound from (4.2) as bases are added according the heuristic in Section 5.3. (b) Average number of necessary Newton iterations as basis elements are added.

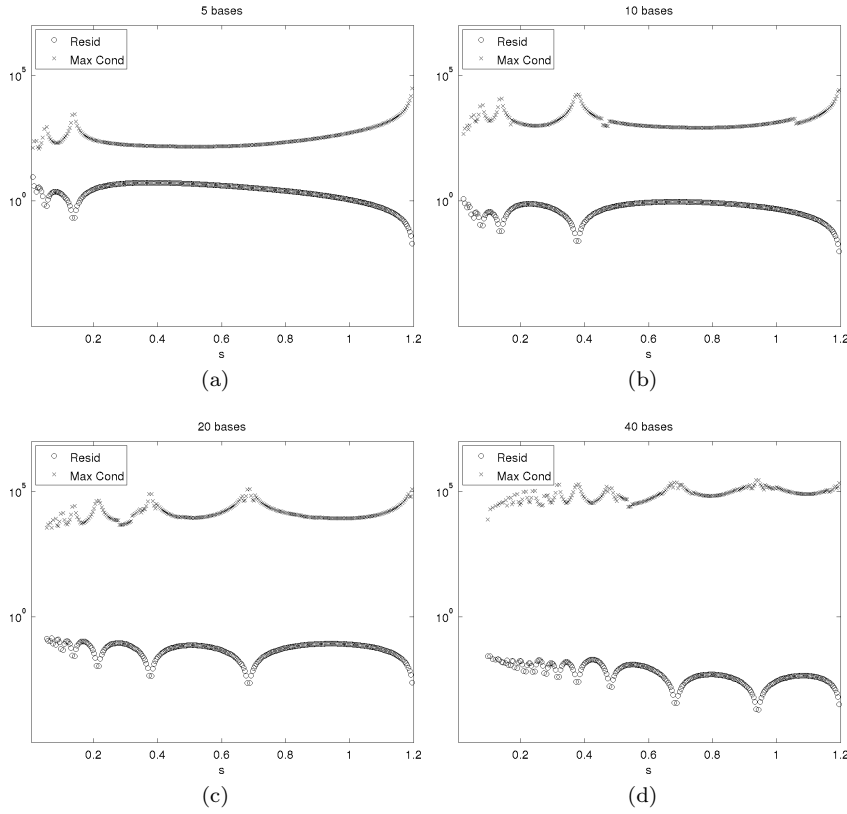


Fig. 6.3: Maximum condition number of R_k from (2.14) over all Newton iterations compared to minimum residual for (a) 5, (b) 10, (c) 20, and (d) 40 bases.

temperature exceeds a critical threshold after 70 seconds. We therefore need the temperature at each point in the domain at $t = 70$ seconds. Given a few simulations of the heat transfer for chosen thermal conductivities, the goal of the model interpolation is to approximate the output of interest at other possible thermal conductivities. We are employing the interpolation at a single point in time, which is the setting where we expect to see computational savings.

6.2.1. Parameterized Model of Thermal Conductivity. In [15], Kodur and co-authors review the high temperature constitutive relationships for steel currently used in European and American standards, as well as present the results of various experimental studies in the research literature. They discuss the challenges of modeling and design given the variation in the standards and experimental data – particularly when designing for safety in high temperature environments. We use their work to inform a statistical model of the thermal conductivity of steel.

To construct a statistical model consistent with the data and codes compiled in [15], we pose a particular model form and choose its parameters to yield good visual agreement with the compiled data. For a given temperature T , let $Y = Y(T, \omega)$ be a random variable that satisfies

$$Y = \bar{Y}(T) + \sigma_Y(T) G_Y(T, \omega). \quad (6.7)$$

The dependence on ω signifies the random component of the model; when clear from the context, we omit explicit dependence on ω . The function $G_Y(T, \omega)$ is a standard Gaussian random field with zero mean and two-point correlation function

$$\mathcal{C}(T_1, T_2) = \exp\left(\frac{-(T_1 - T_2)^2}{\gamma^2}\right). \quad (6.8)$$

The parameter γ controls the correlation between two temperatures T_1 and T_2 in the model. The apparent smoothness of the experimental data (see Figure 6.4) suggests a long correlation length; we choose $\gamma = \sqrt{500}^\circ\text{C}$. The temperature dependent mean $\bar{Y}(T)$ is given by the the log of the Eurocode 3 standard detailed in [15],

$$\bar{Y}(T) = \begin{cases} \log(-0.0333T + 54) & T < 800^\circ\text{C} \\ \log(27.30) & T \geq 800^\circ\text{C} \end{cases} \quad (6.9)$$

The temperature dependent function $\sigma_Y(T)$ is used to scale the variance of the model to be consistent with the experimental data. In particular, we choose

$$\sigma_Y(T) = 0.08 + 0.004\sqrt{T}. \quad (6.10)$$

To model the thermal conductivity $\kappa = \kappa(T, \omega)$, we take the exponential

$$\kappa = \exp(Y), \quad (6.11)$$

which ensures that realizations of κ remain positive. We employ the truncated Karhunen-Loeve expansion [17] of the Gaussian process G_Y to represent the stochasticity in κ as a set of independent parameters:

$$G_Y(T, \omega) \approx \sum_{i=1}^d \phi_i(T) \sqrt{\lambda_i} s_i(\omega), \quad (6.12)$$

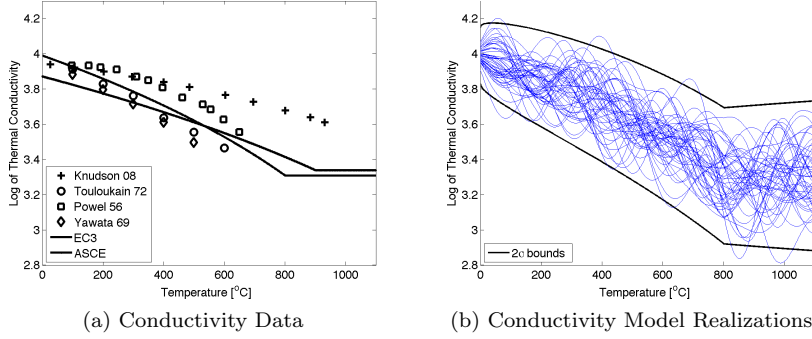


Fig. 6.4: Data and model realizations of temperature dependent thermal conductivity with uncertainty.

where (ϕ_i, λ_i) are eigenpairs of the correlation function \mathcal{C} from (6.8), and $s_i = s_i(\omega)$ are a set of independent standard Gaussian random variables. The eigenfunctions ϕ_i from (6.12) are approximated on a uniform discretization of the temperature interval $[0^\circ\text{C}, 1250^\circ\text{C}]$ with 600 nodes; this discretization is sufficient to capture the correlation effects. To compute the approximate (ϕ_i, λ_i) , we solve the discrete eigenvalue problem for the symmetric, positive semidefinite correlation matrix associated with the discretization of the temperature interval. The exponential decay of the computed eigenvalues justifies a truncation of $d = 11$.

The random variables s_i control the realization of the stochastic model. We can then treat them as a set of independent input parameters for uncertainty and sensitivity analysis. To summarize, we have the following model for thermal conductivity:

$$\kappa = \kappa(T, \xi) = \exp \left[\bar{Y}(T) + \sigma_Y(T) \left(\sum_{i=1}^d \phi_i(T) \sqrt{\lambda_i} s_i \right) \right]. \quad (6.13)$$

Note that we have been loose with the approximation step in (6.12). In the end, we treat the truncated approximation of G_Y to be the true statistical model. In figure 6.4, We plot fifty realizations of the statistical model alongside the log of the data and standards collected in [15].

6.2.2. Heat Conduction Model. Next we incorporate the statistical model for thermal conductivity into a computational simulation. The domain \mathcal{D} is a cross section of a steel beam – shown in Figure 6.5 with labeled boundary segments Γ_1 and Γ_2 . The boundary temperature is prescribed on Γ_1 to increase rapidly up to a maximum value of 1100°C ; this represents rapid heating due to a fire. The boundary segment on Γ_2 is assigned a zero heat flux condition.

We set up the problem with the following heat conduction model. For $x = (x_1, x_2) \in \mathcal{D}$ and $t \in [0, 70]$, let $T = T(x, t, s)$ be the time and space dependent temperature distribution that satisfies the heat conduction model

$$\rho c \frac{\partial T}{\partial t} = -\nabla \cdot (\kappa \nabla T), \quad (6.14)$$

with boundary conditions

$$T = T_b(x, t), \quad x \in \Gamma_1, \quad \text{and} \quad -\kappa \nabla T = 0, \quad x \in \Gamma_2 \quad (6.15)$$

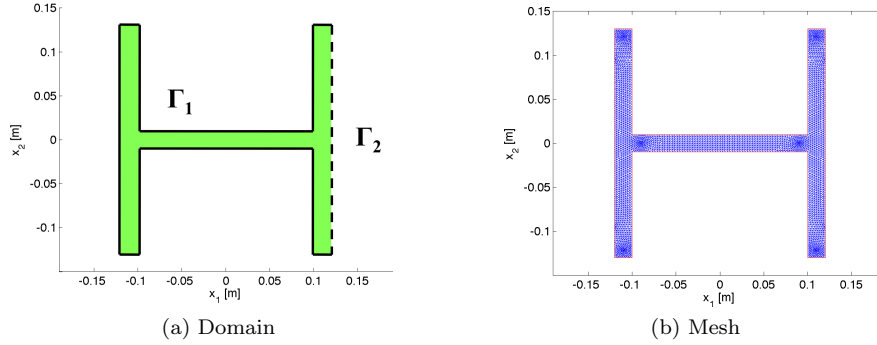


Fig. 6.5: Spatial domain \mathcal{D} for the heat conduction problem (6.14) and associated mesh.

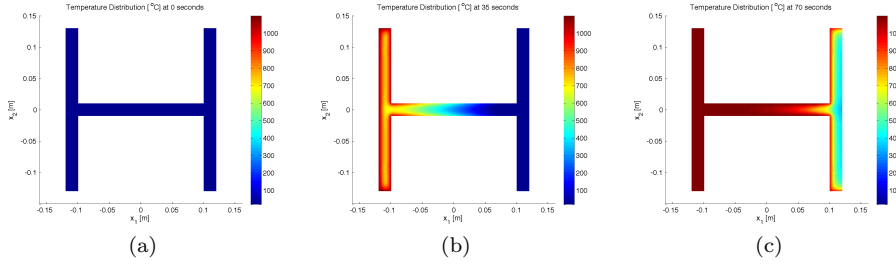


Fig. 6.6: Temperature distribution at the three different times using the mean trend for κ .

where

$$T_b(x, t) = \min \left\{ 1100, \max \left[20, \frac{98}{3}t - 6000x - 700 \right] \right\}. \quad (6.16)$$

The space and time dependent Dirichlet boundary condition represents a rapidly warming environment; its spatial dependence is plotted for three different times in Figure 6.7. Notice that the temperature dependent thermal conductivity κ makes the model nonlinear. We set the initial condition to $T = 20^\circ\text{C}$ throughout the domain.

The spatial domain \mathcal{D} is discretized with 2985 nodes on the irregular triangular mesh shown in Figure 6.5, and the solution is approximated in space with standard piecewise linear finite elements; all spatial discretization is performed with the MATLAB PDE Toolbox. The time stepping is performed by MATLAB's `ode15s` time integrator on the spatially semidiscrete form of (6.14). For reference, the temperature distribution at three points in time is plotted in figure 6.6 using the mean value $\kappa = \exp(\bar{Y})$.

6.2.3. Model Interpolation Study. To test the residual minimizing model interpolation scheme, we set up the following experiment. We first draw 1000 independent realizations of a standard Gaussian random vector with $d = 11$ independent components. Denote these points $s_k \in \mathcal{S}$ with $k = 1, \dots, 1000$ where \mathcal{S} is the input

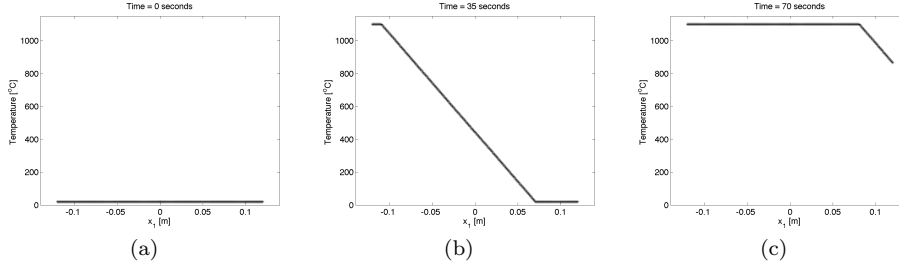


Fig. 6.7: Space and time dependent Dirichlet boundary condition on the boundary segment Γ_1 .

parameter space. For each s_k , we compute the corresponding $\kappa(T, s_k)$, and the resulting temperature distribution $T_k = T(x, t, s_k)$ with the Matlab solver. From each temperature distribution, we compute the fraction of the distribution that exceeds a critical threshold $\tau = 1000^\circ\text{C}$ at time $t = 70$ seconds. Let $Q_k = Q(s_k)$ be defined by the numerical approximation

$$Q_k \approx \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} I(T_k > \tau) dx. \quad (6.17)$$

These Q_k will constitute the cross-validation data set.

To construct the interpolant, we draw 20 independent realizations of an 11-dimensional standard Gaussian random vector; denote these points by $s_j \in \mathcal{S}$. For each s_j with $j = 1, \dots, 20$, we compute $T_j = T(x, t, s_j)$. From each time history of the temperature distribution we retain the distribution at the final time $t = 70$ seconds; these constitute the basis elements for $t = 70$. We also compute the quantity $f(T_j) = -\nabla \cdot (\kappa \nabla T_j)$ at time $t = 70$ for each basis element, which is used to set up the nonlinear least squares problems (2.7).

For each s_k from the cross-validation set, we build the interpolant $\tilde{T}_k = \tilde{T}(x, t = 70, s_k)$ using the 20 basis elements. We then compute the fraction of the interpolated temperature distribution that exceeds the threshold,

$$\tilde{Q}_k \approx \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} I(\tilde{T}_k > \tau) dx. \quad (6.18)$$

We compute the error $\mathcal{E}_k = \mathcal{E}(s_k)$ with respect to the cross-validation data

$$\mathcal{E}_k = |Q_k - \tilde{Q}_k|. \quad (6.19)$$

Errors are averaged over the s_k , $\mathcal{E} = (1/1000) \sum_k \mathcal{E}_k$.

To test the windowing heuristic from Section 5.1, we choose the $M = 5$ basis elements nearest the interpolation point s_k from each basis set. We compute the same \tilde{Q}_k for each s_k using the smaller basis set, and the error is computed as in (6.19).

To average out some of the effects of randomly choosing an especially good or bad basis set with respect to the cross-validation set, we repeat this experiment 10 times with different randomly chosen basis sets. Errors are averaged over all s_k in the cross-validation set and over all 10 experiments. A histogram of the log of all errors

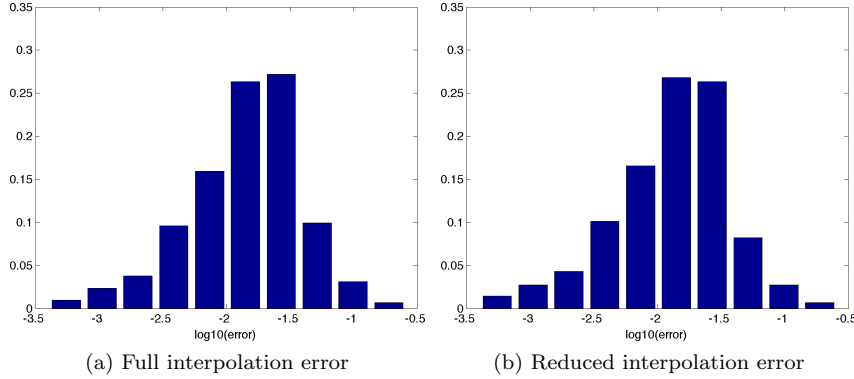


Fig. 6.8: Histograms of the log of the error between the interpolants and the cross-validation data.

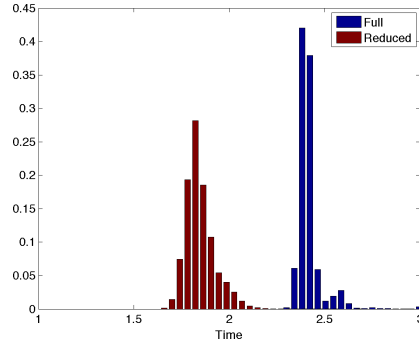


Fig. 6.9: Histogram of the wall clock times for the full and reduced interpolants.

is shown in figure 6.8 for both the full basis set and the windowed basis set. We see that there is practically no difference in error for the smaller windowed basis set.

In figure 6.9, we show a histogram of the wall clock timings for the full basis set and the smaller basis sets. We see that the basis reduction cuts wall clock time by 24% with no practical change in error. In table 6.1, we display the average and standard deviations of the wall clock timings for computing the full model (the cross-validation data), the full interpolation (all 20 bases), and the reduced interpolant (5 chosen bases). We also show the average number of function evaluations in each case; the counts of function evaluations for the full model are output by the Matlab `ode15s` solver. For the full and reduced interpolants, the nonlinear least squares solver used a maximum of ten Newton iterations for each evaluation.

7. Conclusions. We have presented a method for approximating the solution of a parameterized, nonlinear dynamical system using an affine combination of the time histories computed at other input parameter values. The coefficients of the affine combination are computed with a nonlinear least squares procedure that minimizes the residual of the governing equations. In many cases of interest, the computational cost is less than evaluating the full model, which suggests use for reduced order modeling.

Table 6.1: Timing and function evaluation counts for full model and interpolation.

	Avg. Time (s)	Std. Time (s)	Avg. # f Evals
Full Model	178.82	5.59	6166
Full Interp.	2.42	0.09	230
Red. Interp.	1.84	0.08	65

This residual minimizing scheme has similar error and convergence properties to existing reduced basis methods and POD-Galerkin reduced order models; it stands out from existing methods for its ease of implementation by requiring only independent evaluations of the forcing function of the dynamical system. Also, since we do not reduce the basis with a POD type reduction, the approximation interpolates the true time history at the parameter values corresponding to the precomputed solutions.

We proved some interesting properties of this scheme including continuity, convergence, and a lower bound on the error, and we also show that the value of the minimum residual is intimately tied to the conditioning of the least squares problems used to compute the Newton steps. We introduced heuristics to combat this ill-conditioning and further reduce the cost of the method, which we tested with two numerical examples: (i) a three-state dynamical system representing kinetics with a single parameter controlling stiffness and (ii) a nonlinear, parabolic PDE with a high-dimensional random conductivity field.

8. Acknowledgments. The authors thank David Gleich at Purdue University for many helpful conversations and insightful comments. We also thank Venkatesh Kodur at Michigan State University for providing us with the data to motivate and construct the stochastic model of thermal conductivity. We also thank the anonymous reviewers for their helpful and insightful comments.

REFERENCES

- [1] A. C. ANTOUNAS, D. C. SORESENSEN, AND S. GUGERCIN, *A survey of model reduction methods for large-scale systems*, Contemporary Mathematics, 280 (2001), pp. 193–219.
- [2] M. BARRAULT, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, *An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations*, C. R. Acad. Sci. Paris, 339 (2004), pp. 667 – 672.
- [3] V. BARTHELMANN, E. NOVAK, AND K. RITTER, *High dimensional polynomial interpolation on sparse grids*, Advances in Computer Mathematics, 12 (2000), pp. 273–288.
- [4] A. BJORCK, *Numerical Methods for Least Squares Problems*, SIAM, 1996.
- [5] T. BUI-THANH, K. WILLCOX, AND O. GHATTAS, *Model reduction for large-scale systems with high-dimensional parametric input space*, SIAM J. Sci. Comput., 30 (2008), pp. 3270–3288.
- [6] K. CARLBERG, C. BOU-MOSLEH, AND C. FARHAT, *Efficient non-linear model reduction via a least-squares petrov-galerkin projection and compressive tensor approximation*, International Journal for Numerical Methods in Engineering, (2010).
- [7] S. CHATURANTABUT AND D. C. SORESENSEN, *Nonlinear model reduction via discrete empirical interpolation*, SIAM J. Sci. Comput., 32 (2010), pp. 2737 – 2764.
- [8] N. A. C. CRESSIE, *Statistics for Spatial Data*, Wiley, 1993.
- [9] S. C. EISENSTAT AND H. F. WALKER, *Globally convergent inexact newton methods*, SIAM J. Optim., 4 (1994), pp. 393 – 422.
- [10] D. GALBALLY, K. FIDKOWSKI, K. WILLCOX, AND O. GHATTAS, *Non-linear model reduction for uncertainty quantification in large-scale inverse problems*, International Journal for Numerical Methods in Engineering, 81 (2010), pp. 1581 – 1608.
- [11] M. GREPL AND A. T. PATERA, *A posteriori error bounds for reduced-basis approximations of parameterized parabolic partial differential equations*, ESAIM: Mathematical Modelling

- and Numerical Analysis, 39 (2005), pp. 157–181.
- [12] M. C. KENNEDY AND A. O'HAGAN, *Bayesian calibration of computer models*, Journal of the Royal Statistical Society, 63 (2001), pp. 425–464.
 - [13] P. KERFRIDEN, P. GOSSELET, S. ADHIKARI, AND S. BORDAS, *Bridging proper orthogonal decomposition methods and augmented Newton-Krylov algorithms: An adaptive model order reduction for highly nonlinear mechanical problems*, Computer Methods in Applied Mechanics and Engineering, (2010).
 - [14] D. KNOLL AND D. KEYES, *Jacobian-free Newton-Krylov methods: a survey of approaches and applications*, Journal of Computational Physics, 193 (2004), pp. 357–397.
 - [15] V. KODUR, M. DWAIKAT, AND R. FIKE, *High-temperature properties of steel for fire resistance modeling of structures*, Journal of Materials in Civil Engineering, 22 (2010), pp. 423–434.
 - [16] T. G. KOLDA, R. M. LEWIS, AND V. TORCZON, *Optimization by direct search: New perspectives on some classical and modern methods*, SIAM Review, 45 (2003), pp. 385–482.
 - [17] M. LOÈVE, *Probability Theory II*, Springer-Verlag, 1978.
 - [18] N. NGUYEN, A. PATERA, AND J. PERAIRE, *A 'best points' interpolation method for efficient approximation of parameterized functions*, International Journal for Numerical Methods in Engineering, 73 (2008), pp. 521 – 543.
 - [19] N. NGUYEN AND J. PERAIRE, *An efficient reduced-order modeling approach for non-linear parameterized partial differential equations*, International Journal for Numerical Methods in Engineering, 76 (2008), pp. 27 – 55.
 - [20] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer, 2nd ed., 2006.
 - [21] C. PRUD'HOMME, D. V. ROVAS, K. VEROY, L. MACHIELS, Y. MADAY, A. T. PATERA, AND G. TURINICI, *Reliable real-time solution of parametrized partial differential equations: Reduced-basis output bound methods*, Journal of Fluids Engineering, 124 (2002), pp. 70–80.
 - [22] W. H. SCHILDERS, H. A. VAN DER VORST, AND J. ROMMES, eds., *Model Order Reduction: Theory, Research Aspects and Applications*, Springer, 2008.
 - [23] M. L. STEIN, *Interpolation of Spatial Data: Some Theory for Kriging*, Springer, New York, 1999.
 - [24] M. VALORANI, D. A. GAOUSSIS, F. CRETA, AND H. N. NAJM, *Higher order corrections in the approximation of low-dimensional manifolds and the construction of simplified problems with the csp method*, Journal of Computational Physics, 209 (2005), pp. 754 – 786.
 - [25] S. VAN HUFFEL AND S. VANDEWALLE, *The Total Least Squares Problem: Computational Aspects and Analysis*, SIAM, 1991.
 - [26] K. VEROY AND A. T. PATERA, *Certified real-time solution of the parameterized steady incompressible navier-stokes equations: rigorous reduced-basis a posteriori error bounds*, International Journal for Numerical Methods in Fluids, 47 (2005), pp. 77–788.
 - [27] H. WENDLAND, *Scattered Data Approximation*, Cambridge University Press, 2005.
 - [28] D. XIU AND J. S. HESTHAVEN, *High order collocation methods for differential equations with random inputs*, SIAM Journal of Scientific Computing, 27 (2005), pp. 1118 – 1139.